

DZero Offline Computing Update

Qizhong Li

Fermilab


April 27, 2009

All Experimenters' Meeting

D0 Farm Efficiency

Last Fall,

- The efficiency of the D0 farm was less than optimal.
- ~3 weeks behind Data taking.



of jobs slots used (should be saturated at all times).


- CD formed a project to improve D0 farm efficiency.
 - DZero Grid Data Production Initiative
 - Led by Rob Kennedy
 - Included Dzero, CD Rex, system admins and Grid Department.

D0 Grid Production Initiative


- Initiative is a success:
 - Reconfigured the production system SamGrid, which is an interface to Open Science Grid (OSG).
 - Added more forwarding nodes and queuing nodes to farm
 - Separated data production and MC production for using the forwarding nodes
 - Readjusted many parameters in the system
 - Upgraded the batch system
 - Added new nodes to the farm

D0 Farm Efficiency Improvement


- From the end of November 2008, the farm has been running smoothly with high efficiency.
 - The farm job slot usage almost 100%, most of the time.



A Mystery




Farm efficiency improved significantly!



But the event processing rate did not improve?


The Answer: Luminosity Effect



Due to the excellent performance of Tevatron, the average luminosity is increasing

135E30 \rightarrow 175E30.

- The average reconstruction CPU changed:
40 GHz-sec \rightarrow 70 GHz-sec




Increasing CPUs for the farm

- Put CDF retired old nodes onto the data production farm February, 2009.
- After the Moriond 2009 data analysis was over, we temporarily moved some analysis CPUs to the farm to catch up the data production backlog for preparing the data set for Summer 2009 conferences
 - Temporary measure. Would not be feasible in weeks preceding a major conference.
 - Highest backlog was 204 M events behind data taking
 - To expand the farm (moved the analysis nodes to the farm)
1640 CPUs → 3200 CPUs



Production Status

Backlog caught up!



- Expanded the farm:
Mar.12 – Apr.13
2009
- Backlog caught up
- Reconfigured the farm back, return borrowed CPUs to analysis.


Farm Performance



Expanded Farm Job Success Rate


- with the expanded farm, the job failure rates became higher.

- random failures
- resubmit always success
- because all the infrastructure being pushed to its limit.




Event Size

Average Event Size vs Luminosity



Average reconstructed event size





Current Plan

- During 2009 shutdown expect to make half of farm nodes available for analysis for summer conferences.
- The other half of the farm will be devoted to
 - reprocessing of small data sets to address localized issues
 - testing jobs
- Purchase more CPUs for post-shutdown use.

MC Production

- MC production continues to use resources from: OSG, non-grid (In2p3) and native SamGrid sites
- Generate $\sim 24\text{M}$ events / week.



Summary

- D0 offline computing is doing well.
- Data production backlog is caught up.
- We are in preparation for Summer conference data set.
- Need more CPUs long term to keep up with data reconstruction and analysis
- Tested the SamGrid production system scalability, ready for future farm expansion.
- Moved the CPU intensive analysis jobs to grid.
- MC production is progressing well, with substantial grid contribution.
- A lot of people's efforts contributed to success.